

Amendments to the Claims

1 Claim 1 (currently amended): A computer program product for serializing data retrievals and  
2 updates, the computer program product embodied on one or more computer-readable media and  
3 comprising:

4 computer-readable program code means for creating two identical indexes, each  
5 representing an initial state for accessing stored data and each indexing a single copy of data  
6 being indexed;

7 computer-readable program code means for performing searches against the data using a  
8 first of the two indexes;

9 computer-readable program code means for performing a first update against the data  
10 using a second of the two indexes, yielding a revised index;

11 computer-readable program code means for serializing information on how the first  
12 update affected the second index, including how the second index was traversed for making the  
13 first update and how the second index was modified in the first update;

14 computer-readable program code means for switching the first index and the revised  
15 index, responsive to operation of the computer-readable program code means for performing the  
16 first update, such that the first index becomes the second index and the revised index becomes  
17 the first index;

18 computer-readable program code means for applying, after operation of the computer-  
19 readable program code means for switching, the serialized information to the second index, using  
20 the information about how the second index was traversed and modified to efficiently traverse  
21 and modify the newly-switched second index, thereby yielding a second index that is

22 synchronized with, and structurally identical to, the first index; and  
23 computer-readable program code means for performing subsequent searches against the  
24 data using the first index.

1 Claim 2 (previously presented): The computer program product according to Claim 1, further  
2 comprising:

3 computer-readable program code means for obtaining an exclusive lock on the second  
4 index prior to operation of the computer-readable program code means for performing the first  
5 update; and

6 computer-readable program code means for releasing the exclusive lock after operation of  
7 the computer-readable program code means for applying the serialized information.

1 Claim 3 (previously presented): The computer program product according to Claim 1, wherein  
2 atomic instructions are used to maintain proper synchronization between the first index and the  
3 second index.

1 Claim 4 (previously presented): The computer program product according to Claim 1, wherein  
2 the computer-readable program code means for serializing information further comprises  
3 computer-readable program code means for queuing a transaction, and wherein the computer-  
4 readable program code means for applying the serialized information further comprises  
5 computer-readable program code means for applying the queued transaction to the second index.

1 Claim 5 (currently amended): The computer program product according to Claim 1, further  
2 comprising computer-readable program code means for performing a subsequent update against  
3 the data using the second index that results from operation of the computer-readable program  
4 code means for applying the serialized information; and wherein operation of the computer-  
5 readable program code means for performing the subsequent update causes another operation of  
6 the computer-readable program code means for serializing, the computer-readable program code  
7 means for switching, and the computer-readable program code means for applying.

1 Claim 6 (currently amended): A system for serializing data retrievals and updates in a computing  
2 environment, comprising:

3 means for creating two identical indexes, each representing an initial state for accessing  
4 stored data and each indexing a single copy of data being indexed;

5 means for performing searches against the data using a first of the two indexes;

6 means for performing a first update against the data using a second of the two indexes,  
7 yielding a revised index;

8 means for serializing information on how the first update affected the second index,  
9 including how the second index was traversed for making the first update and how the second  
10 index was modified in the first update;

11 means for switching the first index and the revised index, responsive to operation of the  
12 means for performing the first update, such that the first index becomes the second index and the  
13 revised index becomes the first index;

14 means for applying, after operation of the means for switching, the serialized information

15 to the second index, using the information about how the second index was traversed and  
16 modified to efficiently traverse and modify the newly-switched second index, thereby yielding a  
17 second index that is synchronized with, and structurally identical to, the first index; and  
18 means for performing subsequent searches against the data using the first index.

1 Claim 7 (previously presented): The system according to Claim 6, further comprising:  
2 means for obtaining an exclusive lock on the second index prior to operation of the means  
3 for performing the first update; and  
4 means for releasing the exclusive lock after operation of the means for applying the  
5 serialized information.

1 Claim 8 (previously presented): The system according to Claim 6, wherein atomic instructions  
2 are used to maintain proper synchronization between the first index and the second index.

1 Claim 9 (previously presented): The system according to Claim 6, wherein the means for  
2 serializing information further comprises means for queuing a transaction, and wherein the  
3 means for applying the serialized information further comprises means for applying the queued  
4 transaction to the second index.

1 Claim 10 (currently amended): The system according to Claim 6, further comprising means for  
2 performing a subsequent update against the data using the second index that results from  
3 operation of the means for applying the serialized information; and wherein operation of the

4 means for performing the subsequent update causes another operation of the means for  
5 serializing, the means for switching, and the means for applying.

1 Claim 11 (currently amended): A method for serializing data retrievals and updates in a  
2 computing environment, comprising steps of:

3 creating two identical indexes, each representing an initial state for accessing stored data  
4 and each indexing a single copy of data being indexed;

5 performing searches against the data using a first of the two indexes;

6 performing a first update against the data using a second of the two indexes, yielding a  
7 revised index;

8 serializing information on how the first update affected the second index, including how  
9 the second index was traversed for making the first update and how the second index was  
10 modified in the first update;

11 switching the first index and the revised index, responsive to performing the first update,  
12 such that the first index becomes the second index and the revised index becomes the first index;

13 applying, after the switching step, the serialized information to the second index, using  
14 the information about how the second index was traversed and modified to efficiently traverse  
15 and modify the newly-switched second index, thereby yielding a second index that is

16 synchronized with, and structurally identical to, the first index; and

17 performing subsequent searches against the data using the first index.

1 Claim 12 (previously presented): The method according to Claim 11, further comprising steps

2 of:

3 obtaining an exclusive lock on the second index prior to performing the first update; and  
4 releasing the exclusive lock after applying the serialized information.

1 Claim 13 (previously presented): The method according to Claim 11, wherein atomic  
2 instructions are used to maintain proper synchronization between the first index and the second  
3 index.

1 Claim 14 (previously presented): The method according to Claim 11, wherein the step of  
2 serializing information further comprises the step of queuing a transaction, and wherein the step  
3 of applying the serialized information further comprises the step of applying the queued  
4 transaction to the second index.

1 Claim 15 (currently amended): The method according to Claim 11, further comprising the step  
2 of performing a subsequent update against the data using the second index that results from  
3 applying the serialized information; and wherein the step of performing the subsequent update  
4 causes repeating the serializing, switching, and applying steps.

1 Claim 16 (currently amended): A method of serializing access to data in a computing system,  
2 comprising steps of:

3 maintaining two trees as indexes to the data, a first of which is used for searches and a  
4 second of which is used for update operations, each tree having a use count associated therewith;

5 carrying out searches using the search tree, further comprising the steps of:

6 determining, for each new search request, which of the trees is currently the search  
7 tree;

8 incrementing the use count for the search tree;

9 performing the new search request using the search tree; and

10 decrementing the use count for the search tree, responsive to completion of the  
11 performing step; and

12 carrying out each update using the update tree, further comprising the steps of:

13 determining which of the trees is currently the update tree;

14 performing an update against the update tree;

15 serializing, for each update operation, a record of how the update operation  
16 affected the second update tree;

17 switching the two trees update tree to become the search tree and the search tree to  
18 become the update tree, responsive to completion of the steps of performing each the update  
19 operation and serializing the record; and

20 applying the serialized record to the newly-switched second update tree, provided  
21 that the use count for the newly-switched update tree has reached a value that indicates that no  
22 search requests are currently being performed against this newly-switched update tree, delaying  
23 the step of applying the serialized record if necessary until the use count for the newly-switched  
24 update tree has reached this value, and wherein the step of applying the serialized record ensures  
25 such that both the first search tree and the second update tree reflect the each update operation.

1 Claim 17 (previously presented): A method of serializing access to data in a computing system,  
2 comprising steps of:  
3 maintaining two indexes to the data, a first of which is used for searches and a second of  
4 which is used for update operations;  
5 serializing, for each update operation, a record of how the update operation affected the  
6 second index;  
7 switching the two indexes, responsive to performing each update operation; and  
8 applying the serialized record to the newly-switched second index, such that both the first  
9 index and the second index reflect the update operation.

1 Claim 18 (previously presented): The method according to Claim 17, wherein the two indexes  
2 are B-trees.

1 Claim 19 (previously presented): The computer program product according to Claim 1, wherein  
2 the indexes are implemented as trees.

1 Claim 20 (previously presented): The computer program product according to Claim 1, wherein  
2 the indexes are implemented as hash tables.

1 Claim 21 (previously presented): The system according to Claim 6, wherein the indexes are  
2 implemented as trees.



1 Claim 22 (previously presented): The system according to Claim 6, wherein the indexes are  
2 implemented as hash tables.

1 Claim 23 (previously presented): The method according to Claim 11, wherein the indexes are  
2 implemented as trees.

1 Claim 24 (previously presented): The method according to Claim 11, wherein the indexes are  
2 implemented as hash tables.